

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants: Parruck, et al.

Assignee: Azanda Network Devices

Title: "Multi-Service Segmentation and Reassembly Device Involving Multiple Data Path Integrated Circuits" (As Amended)

Serial No.: Unknown

Filed: October 12, 2001

Examiner: Unknown

Group Art Unit: Unknown

Atty. Doc. No.: AZA-003-6D/2001-P009

October 12, 2001

ASSISTANT COMMISSIONER FOR PATENTS
Washington, D.C. 20231

PRELIMINARY AMENDMENT

Sir:

Before action on the merits, and before calculating the filing fee, please amend the above-identified application as follows.

IN THE DRAWINGS:

A minor error in Figure 52 is corrected. **A copy of Figure 52 with the correction indicated in red is attached.** In reviewing the correction, the Examiner's attention is direct to page 59, line 22 of the specification through page 60, line 21 and the corrections made to that text (see below). Review and approval of the correction by the Examiner is requested. No subject matter is added by the correction, nor is it the intent of Applicants or the undersigned that any subject matter be added.

IN THE TITLE:

Please change the title of the application to: "Multi-Service Segmentation and Reassembly Device Involving Multiple Data Path Integrated Circuits".

IN THE SPECIFICATION:

Page 1, before the first sentence of the specification, please add the following header and paragraph:

CROSS-REFERENCE TO RELATED APPLICATION

This application claims priority under 35 U.S.C. §120 from U.S. Patent Application Serial Number 09/851,565, filed May 8, 2001.

The paragraph from page 14, line 31 to page 15, line 10 should be replaced with:

Figure 4 is a simplified diagram of a router 100 in accordance with an embodiment of the present invention. Router 100 includes a plurality of line cards 101-104, a switch fabric 105 and a central processing unit (CPU) 106. The line cards 101-104 are coupled to switch fabric 105 by parallel buses 107-114. In the present example, each of parallel buses 107-114 is a 16-bit SPI-4, Phase II, LVDS parallel bus operating at 400 MHz at a double data rate (DDR). CPU 106 is coupled to line cards 101-104 by another parallel bus 131. In the present example, parallel bus 131 is a 32-bit PCI bus. In this example, each of the line cards can receive network communications in multiple formats. For example, line card 101 is coupled to a fiber optic cable 115 such that line card 101 can receive from cable 115 network communications at OC-192 rates in packets, ATM cells, and/or AAL5 cells. AAL5 cells are considered a type of ATM cell.

Page 33, the first full paragraph, should be replaced with:

In accordance with egress application type 11, segmentation block 203 does not perform segmenting per se but rather forwards the 64-byte chunk to memory manager block 204 via the per-port FIFO mechanism described above. Memory manager block 204 issues an enqueue command via enqueue command line 320, and stores the 64-byte chunk in payload memory 217. Per flow queue block 207 adds the BID to the per flow queue for flow #1. The same process flow occurs for the second and third switch cells associated with flow #1.

Page 34, the second full paragraph, should be replaced with:

Figure 32 illustrates the general flow of information out of egress MS-SAR 200 in one scenario when the linked lists for flow #1 and flow #2 are dequeued. As each BID is dequeued from a per flow queue, the associated 64-byte chunk is read from payload memory 217 and is supplied via 128-bit wide data bus 324 to reassembly block 205. As explained above in connection with the ingress mode, reassembly block 205 maintains one reassembly queue for each of the 64 logical output ports. If two flows share the same logical output port, then the entire linked list for one flow must be dequeued before the linked list for the next flow having the same logical output port is dequeued. In the example of Figure 8, flow #1 and flow #2 have different logical output ports. The dequeue command as received by memory manager block 204 from per flow queue block 207 contains the port ID (PID) of the one of the 64 logical output ports of line card 103.

Page 36, the second full paragraph, should be replaced with:

The AAL5 trailer also contains a sixteen-bit "length of payload" field that indicates the length of the payload of the packet being reassembled. Reassembly block 205 maintains a partial packet byte count value for the packet. As each 64-byte chunk of the packet is received from memory manager block 204, reassembly block 205 adds the number of bytes received to the partial packet byte count value. After the last 64-byte chunk for the packet has been received and the partial packet byte count value has been added to, reassembly block 205 compares the now complete packet byte count value with the "length of payload" value from the AAL5 trailer. The two should match.

Page 59, the fourth full paragraph, should be replaced with:

Figure 52 is a diagram of an external memory device 529 that is coupled to a control integrated circuit. External memory device 529 may, for example, be external memory 225 of Figure 10. External memory device 529 stores two types of information, information #1 and information #2, both of which must be accessed within a particular amount of time related to the rate of incoming information. Where the control integrated circuit is clocked by a clock signal, this particular amount of time can be referred to as

four clock periods. In the example to the left of Figure 52, each piece of information can be accessed in two clock periods. Both pieces of information are stored in the same external memory device requiring one to be accessed before the other. A total of four clock periods is therefore required to access both pieces of information.

Page 60, the second full paragraph, should be replaced with:

In accordance with one novel aspect, two external memories are used. Information #1 is stored in the first external memory 530 and information #2 is stored in second external memory 531. The two external memories are accessed at the same time in parallel. If external memories 530 and 531 are the same type of external memory 529, then these external memories have access times of two clock periods and both information #1 and information #2 are accessed within the required two clock periods. It is to be understood, of course, that the four and two in the example are used only to illustrate the technique of accessing memories in parallel to facilitate handling higher data throughput rates. The technique here is not limited to the numbers in this example. An example of information #1 is cell count information stored in the embodiment of Figure 10 in PFQ STAT memory 225. An example of information #2 is packet count information stored in the embodiment of Figure 10 in PFQ STAT memory 225. These two types of information are, in one embodiment of Figure 48, stored in different external memory devices.

Page 60, line 32 (the header), should be replaced with:

BACKPRESSURING USING SERIAL BUS:

Page 70, the first full paragraph, should be replaced with:

There is at most one segmentation process going on per port. Incoming burst data for a port is accumulated into a corresponding one of the 64-byte buffers in data SRAM 908. The maximum amount of data that can be stored in the 64-byte buffer (either 48, 56, or 64 bytes) is determined by the application type of the flow. The 64-byte format is used for packet data. The 48-byte format is used for AAL5 ATM data. The 56-byte format is used for non-AAL5 ATM data. Control information for the data in a 64-byte

buffer is queued into a location in queue FIFO 912 that corresponds to the particular 64-byte buffer. Segmentation table and statistics information is written back to the segmentation table and statistics memory 907. Input phase state machine 904 contains a six-bit data byte counter (not shown) that increments the six-bit "cell length" (also called the "chunk length") count value (stored in segmentation table memory 907) upon receiving each incoming data byte for a 64-byte buffer. When the number of bytes in the 64-byte buffer reaches the maximum amount allowed by the format used, then the data in the 64-byte buffer is ready for transfer to memory manager block 204.

The paragraph from page 74, line 30 through page 75, line 8, should be replaced with:

Although in this embodiment external memory 217 is ZBT (zero bus turnaround) SRAM, external memory 217 may in other embodiments be another type of memory such as, for example, DRAM. In some embodiments, bandwidth to external memory 217 is increased by realizing external memory 217 in multiple integrated circuit memory devices, where each is accessed via a different one of a plurality of interface ports controlled by the same memory controller block 1007. Memory controller 1007 in Figure 58 is programmable to interface with pipelined memory. In one such embodiment, the four memory locations where a chunk of data is stored are read in sequence. Although the reading of the first of the four locations requires multiple clock cycles, pipelining is employed such that the other of the four locations are read, one location on each subsequent clock cycle.

The paragraph from page 84, line 20 through page 85, line 23, should be replaced with:

Enqueue data pipe block 1115 converts data from 128-bit wide to 64-bit wide for entry into data SRAM 1107. Enqueue data pipe block 1115 also provides the appropriate delay necessary for enqueue state machine block 1101 to obtain the pointer needed to store the chunk into data SRAM 1107. Dequeue data pipe block 1116 delays data to the output FIFO block 1110 so that the header can be inserted and the CRC checked for L2 packets before the last word is sent. Header SRAM block 1106 is a

Applicants: Parruck, et al.
Serial No.: Unknown
Docket No.: AZA-003-6D/2001-P009

256x64 bit internal dual port SRAM that stores header information on a per-port basis. The memory is organized as 128 sixteen-byte buffers so that two sixteen-byte headers are stored for each output port. Output FIFO block 1110 is a 32 x 88 bit FIFO implemented using dual port memory. The output FIFO is large enough to store three maximum size chunks. The output FIFO has two full indications (Almost Full and Full) and two empty indications (Almost Empty and Empty). Almost Full is asserted when the FIFO has ten locations left. Full is asserted when the FIFO has nine or fewer locations left. By indicating that the FIFO is almost full when it has only ten locations left, the dequeue state machine block 1108 can then send one more complete chunk. When the FIFO is one word beyond Almost Full, then the Full signal is asserted thereby indicating that once the current chunk is completed that no more chunks are to be transferred until Full goes inactive. Almost Empty is asserted when one word is left in the output FIFO. Empty is asserted when the output FIFO is empty. Outgoing SPI interface block 206 ORs the two empty signals together to determine when to pop the output FIFO. Outgoing SPI interface block 206 also uses the two empty signals to determine when only one word is left in the FIFO. This is necessary because the minimum POP size is two clock cycles, but when only one word is left the FIFO only has valid data for the first read of the FIFO. If Empty is active but Almost Empty is inactive, then only one word is in the FIFO. Dequeue state machine block 1108 causes data to be loaded into output FIFO block 1110 as long as output port calendar block 1109 indicates chunks are available and output FIFO block 1110 is not full. If output FIFO block 1110 goes full, then the current chunk being transferred is completed before reassembly block 1100 stops sending data to output FIFO block 1110. Outgoing SPI interface block 206 pops data from the output FIFO block 1110 as long as FIFO block 1110 is not empty and outgoing SPI interface block 206 is not inserting control words into the data stream. FIFO block 1110 allows outgoing SPI interface block 206 to control the flow of data from reassembly block 1100 so it has time to add the control words into the output data stream (coming out of outgoing SPI interface block 206) between bursts.

IN THE CLAIMS:

Please amend the claims as follows.

Cancel Claims 1-24 and 29-44, without prejudice.

Add new Claims 45-60 as follows.

45.(New) The system of Claim 25, further comprising:

an aggregation integrated circuit having a first bus interface coupled to the second bus interface of the first integrated circuit, a second bus interface coupled to the second bus interface of the second integrated circuit, and a third bus interface through which the first and second flows of network information pass.

46.(New) The system of Claim 45, wherein the first and second flows of network information pass from the third bus interface of the aggregation integrated circuit to the switch fabric.

47.(New) The system of Claim 45, wherein the first and second flows of network information pass from the third bus interface of the aggregation integrated circuit to a framer/mapper.

48.(New) The system of Claim 25, further comprising:

a distribution integrated circuit having a first bus interface coupled to the first bus interface of the first integrated circuit, a second bus interface coupled to the first bus interface of the second integrated circuit, and a third bus interface into which the first and second flows of network information pass.

49.(New) The system of Claim 48, wherein the first and second flows of network information pass from a classification engine and into the third bus interface of the distribution integrated circuit.

Applicants: Parruck, et al.
Serial No.: Unknown
Docket No.: AZA-003-6D/2001-P009

50.(New) The system of Claim 48, wherein the first and second flows of network information pass from the switch fabric and into the third bus interface of the distribution integrated circuit.

51.(New) The system of Claim 48, wherein the distribution integrated circuit receives a packet via the third bus interface of the distribution integrated circuit, wherein the distribution integrated circuit adds a sequence number to the packet, and wherein the distribution integrated circuit outputs the packet with the sequence number to a selected one of the first integrated circuit and the second integrated circuit.

52.(New) The system of Claim 48, wherein the distribution integrated circuit receives a flow of packets on the third bus interface, wherein the distribution integrated circuit adds a sequence number to each packet in the flow such that each respective packet of the flow has a sequence number that is greater than the sequence number of the previous packet in the flow, the distribution integrated circuit outputting the flow of packets having the sequence numbers.

53.(New) A system, comprising:

a distribution integrated circuit having a first bus interface, a second bus interface, and a third bus interface, the distribution integrated circuit receiving a plurality of flows of packets via the third bus interface, the distribution integrated circuit adding a sequence number to each of the packets of a flow such that each successive packet of a flow carries a sequence number that is greater than the sequence number of the previous packet in that flow;

a first integrated circuit having a first bus interface and a second bus interface, a data path extending from the first bus interface of the first integrated circuit, through segmentation circuitry on the first integrated circuit, through reassembly circuitry on the first integrated circuit, and to the second bus interface of the first integrated circuit, the first integrated circuit receiving a first packet containing a first sequence number from the distribution integrated circuit via the first bus interface of the first integrated circuit;

a second integrated circuit having a first bus interface and a second bus interface, the second integrated circuit being substantially structurally identical to the first integrated circuit, a data path extending from the first bus interface of the second integrated circuit, through segmentation circuitry on the second integrated circuit, through reassembly circuitry on the second integrated circuit, and to the second bus interface of the second integrated circuit, the second integrated circuit receiving a second packet containing a second sequence number from the distribution integrated circuit via the first bus interface of the second integrated circuit; and

a control integrated circuit coupled to the first integrated circuit and to the second integrated circuit, the first integrated circuit supplying the first sequence number to the control integrated circuit, the second integrated circuit supplying the second sequence number to the control integrated circuit, the control integrated circuit using the first and second sequence numbers to determine which packets of which flows have been received onto which of the first and second integrated circuits.

54.(New) The system of Claim 53, wherein the first integrated circuit performs a lookup operation on the first packet and identifies a flow identifier associated with the first packet, the first integrated circuit supplying the flow identifier to the control integrated circuit, the control integrated circuit using both the flow identifier and the first sequence number in said determining of which packets of which flows have been received onto which of the first and second integrated circuits.

55.(New) The system of Claim 53, wherein the control integrated circuit maintains a packet queue for each flow of the plurality of flows.

56.(New) The system of Claim 53, wherein the control integrated circuit controls the first integrated circuit and the second integrated circuit such that the system performs traffic shaping and traffic metering.

57.(New) The system of Claim 53, further comprising:

an aggregation integrated circuit that receives packets from the first and second integrated circuits and that outputs the packets, the packets received being packets of a particular flow, the packets of the particular flow having been received by the distribution integrated circuit in a particular order, wherein the aggregation integrated circuit outputs the packets of the particular flow from the aggregation integrated circuit in the same order in which the packets of the particular flow were received by the distribution integrated circuit.

58.(New) A method, comprising:

receiving onto a first integrated circuit a flow of packets, the packets of the flow being received in a particular order, and distributing the packets from the first integrated circuit to a plurality of data path integrated circuits;

processing the packets on the plurality of data path integrated circuits, each of the data path integrated circuits having segmentation and reassembly circuitry, each of the plurality of data path integrated circuits being substantially structurally identical to every other of the plurality of data path integrated circuits; and

receiving onto a second integrated circuit the packets from the data path integrated circuits and aggregating the packets such that the packets are output from the second integrated circuit in the same order that they were received onto the first integrated circuit, wherein the first integrated circuit also receives a flow of cells, the cells being received in a particular order, the cells being processed by the plurality of data integrated circuits, the second integrated circuit outputting the cells in the same order that they were received onto the first integrated circuit, and wherein the first integrated circuit and the plurality of data path integrated circuits and the second integrated circuit are all disposed on a single line card.

59.(New) A system, comprising:

an aggregation integrated circuit; and

means, coupled to the aggregation integrated circuit, for receiving a flow of packets, the packets being received in a particular order, the means also being for

Applicants: Parruck, et al.
Serial No.: Unknown
Docket No.: AZA-003-6D/2001-P009

processing the packets on a first data path integrated circuit and on a second data path integrated circuit, some of the packets being processed on the first data path integrated circuit, others of the packets being processed on the second data path integrated circuit, each of the first and second data path integrated circuits having segmentation circuitry and reassembly circuitry, the packets being supplied from the first and second data path integrated circuits to the aggregation integrated circuit such that the packets are output from the aggregation integrated circuit in the same particular order in which the packets were received by the means.

60.(New) The system of Claim 59, wherein the means comprises the first data path integrated circuit, the second data path integrated circuit, a distribution integrated circuit, and a control integrated circuit, the control integrated circuit being coupled to the first and second data path integrated circuits, the distribution integrated circuit receiving the packets and distributing the packets to the first and second data path integrated circuits.

REMARKS

Consideration and allowance is respectfully requested. Before examination, please make the amendments as set forth above. After entry of this Preliminary Amendment, Claims 25-28 and 45-60 are pending. If the Examiner would like to discuss any aspect of this application, the Examiner is requested to contact the undersigned at (925) 485-9923.

I hereby certify that this correspondence is being deposited with the United States Postal Service as "Express Mail Post Office to Addressee" addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231, on

October 12, 2001, as

Express Mail No.: EL928548105US.

T. Lester Wallace

Lester Wallace
Signature

Oct. 12, 2001
Date of Signature

Respectfully submitted,

Lester Wallace

T. Lester Wallace
Attorney for Applicants
Reg. No. 34,748

VERSION WITH MARKINGS TO SHOW CHANGES MADE

THE DRAWINGS:

A copy of Figure 52 with the corrections indicated in red is attached.

THE SPECIFICATION:

Page 1, before the first sentence of the specification, please add the following header and paragraph:

CROSS-REFERENCE TO RELATED APPLICATION

This application is a divisional of application serial number 09/851,565, filed May 8, 2001.

The paragraph from page 14, line 31 through page 15, line 10:

Figure 4 is a simplified diagram of a router 100 in accordance with an embodiment of the present invention. Router 100 includes a plurality of line cards 101-104, a switch fabric 105 and a central processing unit (CPU) 106. The line cards 101-104 are coupled to switch fabric 105 by parallel buses 107-114. In the present example, each of parallel buses 107-114 is a 16-bit SPI-4, Phase II, LVDS parallel bus operating at 400 MHz at a double data rate (DDR). CPU 106 is coupled to line cards 101-104 by another parallel bus 131. In the present example, parallel bus [130] **131** is a 32-bit PCI bus. In this example, each of the line cards can receive network communications in multiple formats. For example, line card 101 is coupled to a fiber optic cable 115 such that line card 101 can receive from cable 115 network communications at OC-192 rates in packets, ATM cells, and/or AAL5 cells. AAL5 cells are considered a type of ATM cell.

Page 33, the first full paragraph:

In accordance with egress application type 11, segmentation block 203 does not **perform** segmenting per se but rather forwards the 64-byte chunk to memory manager block 204 via the per-port FIFO mechanism described above. Memory manager block 204 issues an enqueue command via enqueue command line 320, and stores the 64-byte chunk in payload memory 217. Per flow queue block 207 adds the BID to the per flow queue for flow #1. The same process flow occurs for the second and third switch cells associated with flow #1.

Page 34, the second full paragraph:

Figure 32 illustrates the general flow of information out of egress MS-SAR 200 in one scenario when the linked lists for flow #1 and flow #2 are dequeued. As each BID is dequeued from a per flow queue, the associated 64-byte chunk is read from payload memory 217 and is supplied via [128-bite] **128-bit** wide data bus 324 to reassembly block 205. As explained above in connection with the ingress mode, reassembly block 205 maintains one reassembly queue for each of the 64 logical output ports. If two flows share the same logical output port, then the entire linked list for one flow must be dequeued before the linked list for the next flow having the same logical output port is dequeued. In the example of Figure 8, flow #1 and flow #2 have different logical output ports. The dequeue command as received by memory manager block 204 from per flow queue block 207 contains the port ID (PID) of the one of the 64 logical output ports of line card 103.

Page 36, the second full paragraph:

The AAL5 trailer also contains a sixteen-bit “length of payload” field that indicates the length of the payload of the packet being reassembled. Reassembly block 205 maintains a partial packet byte count value for the packet. As each 64-byte chunk of the packet is received from memory manager block 204, reassembly block 205 adds the number of bytes received to the partial packet byte count value. After the last 64-byte chunk for the packet has been received and the partial packet byte count value has been added to, reassembly block 205 compares the now complete packet byte count value with the “length of payload” value from the AAL5 trailer. The two should match.

Page 59, the fourth full paragraph:

Figure 52 is a diagram of an external memory device 529 that is coupled to a control integrated circuit. External memory device 529 may, for example, be external memory 225 of Figure 10. External memory device 529 stores two types of information, information #1 and information #2, both of which must be accessed within a particular amount of time related to the rate of incoming information. Where the control integrated circuit is clocked by a clock signal, this particular amount of time can be referred to as

[eight] **four** clock periods. In the example to the left of Figure 52, each piece of information can be accessed in two clock periods. Both pieces of information are stored in the same external memory device requiring one to be accessed before the other. A total of four clock periods is therefore required to access both pieces of information.

Page 60, the second full paragraph:

In accordance with one novel aspect, two external memories are used. Information #1 is stored in the first external memory 530 and information #2 is stored in second external memory 531. The two external memories are accessed at the same time in parallel. If external memories 530 and 531 are the same type of external memory 529, then these external memories have access times of two clock periods and both information #1 and information #2 are accessed within the required two clock periods. It is to be understood, of course, that the [eight] **four** and two in the example are used only to illustrate the technique of accessing memories in parallel to facilitate handling higher data throughput rates. The technique here is not limited to the numbers in this example. An example of information #1 is cell count information stored in the embodiment of Figure 10 in PFQ STAT memory 225. An example of information #2 is packet count information stored in the embodiment of Figure 10 in PFQ STAT memory 225. These two types of information are, in one embodiment of Figure 48, stored in different external memory devices.

Page 60, line 32 (the header):

[BACKPRESSING] **BACKPRESSURING** USING SERIAL BUS:

Page 70, the first full paragraph:

There is at most one segmentation process going on per port. Incoming burst data for a port is accumulated into a corresponding one of the 64-byte buffers in data SRAM 908. The maximum amount of data that can be stored in the 64-byte [bluffer] **buffer** (either 48, 56, or 64 bytes) is determined by the application type of the flow. The 64-byte format is used for packet data. The 48-byte format is used for **AAL5** ATM data. The 56-byte format is used for [other] **non-AAL5** ATM data. Control information for the data in a 64-byte buffer is queued into a location in queue FIFO 912 that corresponds to

the particular 64-byte buffer. Segmentation table and statistics information is written back to the segmentation table and statistics memory 907. Input phase state machine 904 contains a six-bit data byte counter (not shown) that increments the six-bit “cell length” (also called the “chunk length”) count value (stored in segmentation table memory 907) upon receiving each incoming data byte for a 64-byte buffer. When the number of bytes in the 64-byte buffer reaches the maximum amount allowed by the format used, then the data in the 64-byte buffer is ready for transfer to memory manager block 204.

The paragraph from page 74, line 30 through page 75, line 8:

Although in this embodiment external memory 217 is ZBT (zero bus turnaround) SRAM, external memory 217 may in other embodiments be another type of memory such as, for example, DRAM. In some embodiments, bandwidth to external memory 217 is increased by realizing external memory 217 in multiple integrated circuit memory devices, where each is accessed via a different one of a plurality of interface ports controlled by the same memory controller block 1007. Memory controller 1007 in Figure 58 is programmable to interface with pipelined memory. In one such embodiment, the four memory locations where a chunk of data is stored are read in sequence. Although the reading of the first of the four locations requires multiple clock cycles, pipelining is employed such that the other of the four locations are read, one location on each subsequent clock cycle.

The paragraph from page 84, line 20 through page 85, line 23:

Enqueue data pipe block 1115 converts data from 128-bit wide to 64-bit wide for entry into data SRAM 1107. Enqueue data pipe block 1115 also provides the appropriate delay necessary for enqueue state machine block 1101 to obtain the pointer needed to store the chunk into data SRAM 1107. Dequeue data pipe block 1116 delays data to the output FIFO block 1110 so that the header can be inserted and the CRC checked for L2 packets before the last word is sent. Header SRAM block 1106 is a 256x64 bit internal dual port SRAM that stores header information on a per-port basis. The memory is organized as 128 sixteen-byte buffers so that two sixteen-byte headers are stored for each output port. Output FIFO block 1110 is a 32 x 88 bit FIFO

implemented using dual port memory. The [Th e] output FIFO is large enough to store three maximum size chunks. The output FIFO has two full indications (Almost Full and Full) and two empty indications (Almost Empty and Empty). Almost Full is asserted when the FIFO has ten locations left. Full is asserted when the FIFO has nine or fewer locations left. By indicating that the FIFO is almost full when it has only ten locations left, the dequeue state machine block 1108 can then send one more complete chunk. When the FIFO is one word beyond Almost Full, then the Full signal is asserted thereby indicating that once the current chunk is completed that no more chunks are to be transferred until Full goes inactive. Almost Empty is asserted when one word is left in the output FIFO. Empty is asserted when the output FIFO is empty. Outgoing SPI interface block 206 ORs the two empty signals together to determine when to pop the output FIFO. Outgoing SPI interface block 206 also uses the two empty signals to determine when only one word is left in the FIFO. This is necessary because the minimum POP size is two clock cycles, but when only one word is left the FIFO only has valid data for the first read of the FIFO. If Empty is active but Almost Empty is inactive, then only one word is in the FIFO. Dequeue state machine block 1108 causes data to be loaded into output FIFO block 1110 as long as output port calendar block 1109 indicates chunks are available and output FIFO block 1110 is not full. If output FIFO block 1110 goes full, then the current chunk being transferred is completed before reassembly block 1100 stops sending data to output FIFO block 1110. Outgoing SPI interface block 206 pops data from the output FIFO block 1110 as long as FIFO block 1110 is not empty and outgoing SPI interface block 206 is not inserting control words into the data stream. FIFO block 1110 allows outgoing SPI interface block 206 to control the flow of data from reassembly block 1100 so it has time to add the control words into the output data stream (coming out of outgoing SPI interface block 206) between bursts.

THE CLAIMS:

Cancel Claims 1-24 and 29-44.

Add new Claims 45-60 as follows.

45.(New) The system of Claim 25, further comprising:

an aggregation integrated circuit having a first bus interface coupled to the second bus interface of the first integrated circuit, a second bus interface coupled to the second bus interface of the second integrated circuit, and a third bus interface through which the first and second flows of network information pass.

46.(New) The system of Claim 45, wherein the first and second flows of network information pass from the third bus interface of the aggregation integrated circuit to the switch fabric.

47.(New) The system of Claim 45, wherein the first and second flows of network information pass from the third bus interface of the aggregation integrated circuit to a framer/mapper.

48.(New) The system of Claim 25, further comprising:

a distribution integrated circuit having a first bus interface coupled to the first bus interface of the first integrated circuit, a second bus interface coupled to the first bus interface of the second integrated circuit, and a third bus interface into which the first and second flows of network information pass.

49.(New) The system of Claim 48, wherein the first and second flows of network information pass from a classification engine and into the third bus interface of the distribution integrated circuit.

50.(New) The system of Claim 48, wherein the first and second flows of network information pass from the switch fabric and into the third bus interface of the distribution integrated circuit.

51.(New) The system of Claim 48, wherein the distribution integrated circuit receives a packet via the third bus interface of the distribution integrated circuit, wherein the distribution integrated circuit adds a sequence number to the packet, and wherein the

distribution integrated circuit outputs the packet with the sequence number to a selected on one of the first integrated circuit and the second integrated circuit.

52.(New) The system of Claim 48, wherein the distribution integrated circuit receives a flow of packets on the third bus interface, wherein the distribution integrated circuit adds a sequence number to each packet in the flow such that each respective packet of the flow has a sequence number that is greater than the sequence number of the previous packet in the flow, the distribution integrated circuit outputting the flow of packets having the sequence numbers.

53.(New) A system, comprising:

a distribution integrated circuit having a first bus interface, a second bus interface, and a third bus interface, the distribution integrated circuit receiving a plurality of flows of packets via the third bus interface, the distribution integrated circuit adding a sequence number to each of the packets of a flow such that each successive packet of a flow carries a sequence number that is greater than the sequence number of the previous packet in that flow;

a first integrated circuit having a first bus interface and a second bus interface, a data path extending from the first bus interface of the first integrated circuit, through segmentation circuitry on the first integrated circuit, through reassembly circuitry on the first integrated circuit, and to the second bus interface of the first integrated circuit, the first integrated circuit receiving a first packet containing a first sequence number from the distribution integrated circuit via the first bus interface of the first integrated circuit;

a second integrated circuit having a first bus interface and a second bus interface, the second integrated circuit being substantially structurally identical to the first integrated circuit, a data path extending from the first bus interface of the second integrated circuit, through segmentation circuitry on the second integrated circuit, through reassembly circuitry on the second integrated circuit, and to the second bus interface of the second integrated circuit, the second integrated circuit receiving a second packet containing a second sequence number from the distribution integrated circuit via the first bus interface of the second integrated circuit; and

a control integrated circuit coupled to the first integrated circuit and to the second integrated circuit, the first integrated circuit supplying the first sequence number to the control integrated circuit, the second integrated circuit supplying the second sequence number to the control integrated circuit, the control integrated circuit using the first and second sequence numbers to determine which packets of which flows have been received onto which of the first and second integrated circuits.

54.(New) The system of Claim 53, wherein the first integrated circuit performs a lookup operation on the first packet and identifies a flow identifier associated with the first packet, the first integrated circuit supplying the flow identifier to the control integrated circuit, the control integrated circuit using both the flow identifier and the first sequence number in said determining of which packets of which flows have been received onto which of the first and second integrated circuits.

55.(New) The system of Claim 53, wherein the control integrated circuit maintains a packet queue for each flow of the plurality of flows.

56.(New) The system of Claim 53, wherein the control integrated circuit controls the first integrated circuit and the second integrated circuit such that the system performs traffic shaping and traffic metering.

57.(New) The system of Claim 53, further comprising:

an aggregation integrated circuit that receives packets from the first and second integrated circuits and that outputs the packets, the packets received being packets of a particular flow, the packets of the particular flow having been received by the distribution integrated circuit in a particular order, wherein the aggregation integrated circuit outputs the packets of the particular flow from the aggregation integrated circuit in the same order in which the packets of the particular flow were received by the distribution integrated circuit.

58.(New) A method, comprising:

receiving onto a first integrated circuit a flow of packets, the packets of the flow being received in a particular order, and distributing the packets from the first integrated circuit to a plurality of data path integrated circuits;

processing the packets on the plurality of data path integrated circuits, each of the data path integrated circuits having segmentation and reassembly circuitry, each of the plurality of data path integrated circuits being substantially structurally identical to every other of the plurality of data path integrated circuits; and

receiving onto a second integrated circuit the packets from the data path integrated circuits and aggregating the packets such that the packets are output from the second integrated circuit in the same order that they were received onto the first integrated circuit, wherein the first integrated circuit also receives a flow of cells, the cells being received in a particular order, the cells being processed by the plurality of data integrated circuits, the second integrated circuit outputting the cells in the same order that they were received onto the first integrated circuit, and wherein the first integrated circuit and the plurality of data path integrated circuits and the second integrated circuit are all disposed on a single line card.

59.(New) A system, comprising:

an aggregation integrated circuit; and

means, coupled to the aggregation integrated circuit, for receiving a flow of packets, the packets being received in a particular order, the means also being for processing the packets on a first data path integrated circuit and on a second data path integrated circuit, some of the packets being processed on the first data path integrated circuit, others of the packets being processed on the second data path integrated circuit, each of the first and second data path integrated circuits having segmentation circuitry and reassembly circuitry, the packets being supplied from the first and second data path integrated circuits to the aggregation integrated circuit such that the packets are output from the aggregation integrated circuit in the same particular order in which the packets were received by the means.

60.(New) The system of Claim 59, wherein the means comprises the first data path integrated circuit, the second data path integrated circuit, a distribution integrated circuit, and a control integrated circuit, the control integrated circuit being coupled to the first and second data path integrated circuits, the distribution integrated circuit receiving the packets and distributing the packets to the first and second data path integrated circuits.






FIG. 49



PACKET QUEUE

FIG. 50



PACKET QUEUE

FIG. 51

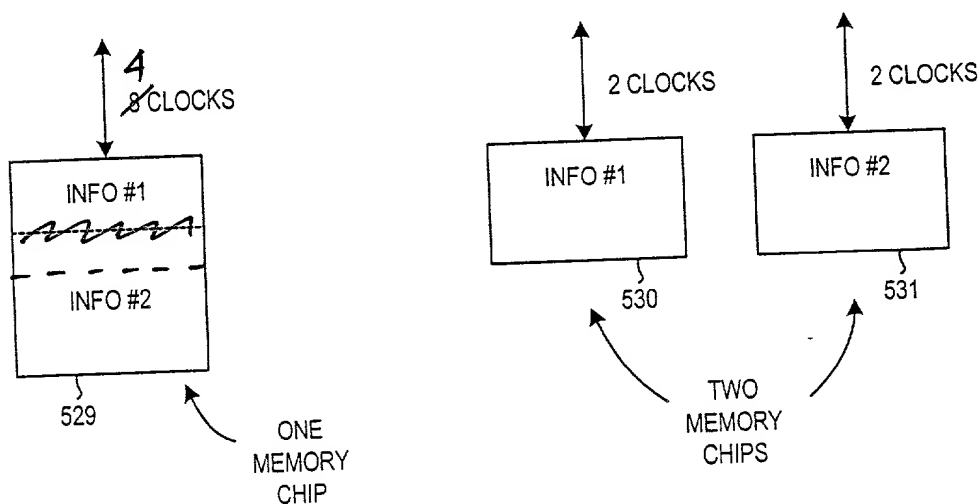


FIG. 52

FIG. 49